Proceedings of the 1998 Winter Simulation Conference
D.J. Medeiros, E.F. Watson, J.S. Carson and M.S. Manivannan, eds.

AN APPROACH TO RANKING AND SELECTION FOR MULTIPLE PERFORMANCE MEASURES

Douglas J. Morrice
John Butler

Department of Management Science and Information Systems
CBA 5.202
The University of Texas at Austin
Austin, Texas 78712-1175, U.S.A.

Peter W. Mullarkey

End-to-End Simulation Department
8311 North RR 620
Schlumberger Austin Product Center
Austin, Texas 78726, U.S.A.

ABSTRACT

In this paper, we develop a ranking and selection procedure for making multiple comparisons of systems that have multiple performance measures. The procedure combines multiple attribute utility theory with ranking and selection to select the best configuration from a set of K configurations using the indifference zone approach. We demonstrate our procedure on a simulation model of a large project that has six performance measures.

1 INTRODUCTION

In recent work, Morrice et al. (1997) developed a simulation model of a project that contains multiple input parameters and multiple performance measures. This paper details how we use the simulation and ranking and selection (R&S) to select the best project configuration over K (>1) possible configurations. The K configurations are constructed from different settings of the input parameters.

Evaluating project configurations on multiple performance measures complicates the R&S analysis. Most of the R&S literature focuses on procedures that are designed for scalar performance measures (see, for example, Bechhofer, Santner, and Goldsman 1995). There are at least three ways to deal with this problem. The first approach is to extend the theory and develop multiple variable R&S procedures. In a business setting, a second approach is often used: convert project performance over multiple measures to a scalar measure using costs. Costing has many obvious advantages but it has some disadvantages, as well. For example, accurate cost data may not be available due to insufficient resources. Additionally, it may be impossible to accurately cost intangibles (e.g., the quality of life, good will, etc.) even if the resources are available.

A third approach is to convert multiple performance measures to a scalar performance measure using multiple attribute utility (MAU) theory. MAU theory can be used instead of a costing approach when good cost data are not available. Alternatively, MAU theory can be used to embellish costing information that is considered to be incomplete (e.g., to account for the intangibles).

In this paper, we focus on the third approach and combine multiple attribute utility theory with statistical R&S using the indifference zone approach. The goal is to select the best project configuration from a set of K configurations when project performance is measured over multiple performance measures.

The remainder of the paper is organized in the following manner. Section 2 describes the project example from Morrice et al. (1997) that will be used throughout the paper. Section 3 contains a brief overview of the MAU theory. Section 4 provides the set-up for R&S and a description of the combined R&S and MAU procedure. Section 5 discusses one of the main research issues: the selection of the indifference zone parameter δ^*. Section 6 illustrates application of the procedure on the project example described in Section 2. Section 7 contains some concluding remarks.

2 EXAMPLE

We use the methodology developed in this paper to analyze the simulation model of the project described in Morrice et al. (1997). The simulation models a large outdoor operation called a signal quality survey. Signal quality surveys are conducted over large geographical areas (tens to hundreds of square kilometers). They are projects taking anywhere from a few days to a few years with the number of personnel ranging from 20 to 1000 people, requiring capital equipment valued in the tens of millions of dollars, and generating survey revenues ranging from hundreds of thousands to hundreds of millions of dollars. The simulation model was designed to support bidding, planning, and conducting these large, complicated, and expensive projects in a profitable manner.

The execution of a signal quality survey requires the coordination of five types of crews (see Figure 1). Briefly, the signal crew sends signals from several geographic locations that are recorded by the recording crew. The layout crew places receiving (or monitoring) equipment at several geographic locations so that the recording crew can receive signals sent by the source crew. The transport crew brings the layout crew receiving equipment. The packing crew prepares receiving equipment for the transport crew that is no longer required on a particular part of a survey for receiving signals sent by the signal crew.

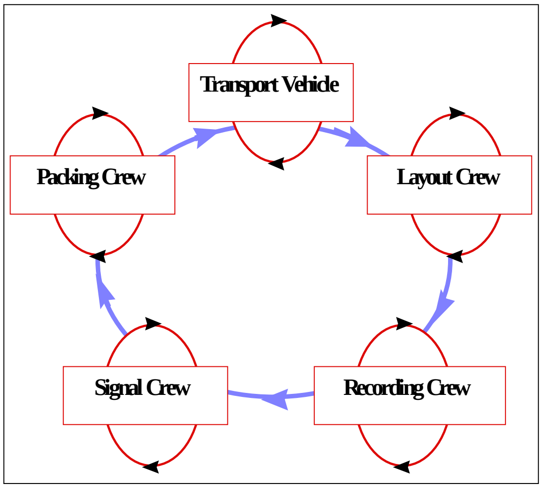

Figure 1: Crews in a Signal Quality Survey

Performance measures on this project include percent utilization for all crew types, project duration, and cost. We will model four project configurations differentiated by the number of source crews and the amount of receiving equipment available. Each configuration will be evaluated on the multiple performance measures and the best will be selected (see Section 6).

3 MAU THEORY: AN OVERVIEW

MAU theory (Keeney and Raiffa, 1976) is one of the major analytical tools associated with the field of decision analysis (see, for example, Clemen 1991). Simply, decision analysis is a logical and formal approach to the solution of problems that are too complex to solve informally. In the past, decision analysis has been applied to problems such as siting an electricity generation facility (Keeney, 1980), choosing among vendors for the evaluation of alternatives for the commercial generation of electricity by nuclear fusion (Dyer and Lorber, 1982), and selecting a nuclear waste clean up strategy (Keeney and von Winterfeldt, 1994).

A MAU analysis of alternatives (in our example, project configurations) explicitly identifies the measures that are used to evaluate the alternatives, and helps to identify those alternatives that perform well on a majority of these measures, with a special emphasis on the measures that are considered to be relatively more important. In order to carry out the analysis, some facts regarding each of the alternatives are required, and in some cases some assumptions will be needed to estimate the performance of the alternatives on the measures. As an example, different assumptions may lead to optimistic and pessimistic cost estimates for the alternatives.

The MAU methodology for the evaluation of a set of alternatives typically consists of the following steps:

1. Identification of alternatives and measures,
2. Estimation of the performance of the alternatives with respect to the measures,
3. Development of utilities and weights for the measures, and
4. Evaluation of the alternatives and sensitivity analysis.

The alternatives and the measures form a matrix in which each row corresponds to an alternative and each column represents a measure. The cells of the matrix contain estimates of the performance of each alternative on each of the measures. When these estimates are uncertain, it is often appropriate to quantify them with ranges or with probability distributions determined using risk analysis methods, i.e., simulation (e.g., Clemen, 1991; Keeney and von Winterfeldt, 1991).

Step three generates a single attribute utility function over each measure that is scaled from 0 to 1, a weight for each measure, and a multiple attribute utility function derived from the single attribute utility functions and the weights. A single attribute utility function is a scoring function that maps a performance measure from its range of possible values to [0,1]. Common forms of this function include concave for risk averse behavior, convex for risk seeking behavior, linear for risk neutral behavior, and "S" shaped for a hybrid of the convex and concave functions. For theoretical and practical reasons, one popular form for single attribute utility function is

$$U(X) = A - Be^{-(X/RT)} \qquad (1)$$

(Clemen, 1991, page 379). The quantities A, B, and RT are parameters that must be set by the decision maker. Several assessment techniques exist for eliciting utility functions from decision makers, i.e., for setting the parameters A, B and RT in the case of (1) (Logical Decisions, 1996, page 113). Figure 2 contains a graph of (1) for the productivity utility of transport vehicle utilization in our project example where A, B, and RT are approximately equal to 1.019, 2.679, and 0.2, respectively. See Section 6 for additional information on this utility function.

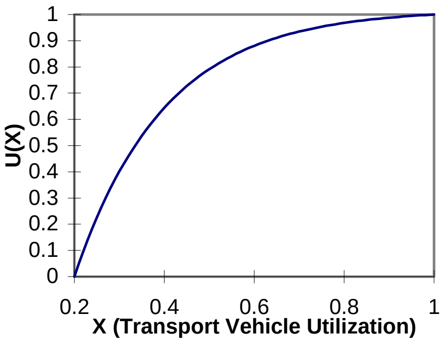

Figure 2: Utility Function for Transport Vehicle Utilization

Several methods also exist for assigning weights to the performance measures (Schoemaker and Waid, 1982, Logical Decisions, 1996, page 130). For example, a method called the Trade-off method, includes all *n* (> 1) performance measures in *n-1* pairwise tradeoffs. In each tradeoff the decision maker is asked to judge on which measure it is more important to improve performance. This procedure in conjunction with the constraint that the weights must sum to one uniquely determines weights. Another popular method is the Analytical Hierarchy Process (AHP) by Saaty (1988).

3.1 Aggregation with Multiple Attribute Utility Functions

Once the performance of each alternative on each measure in the alternatives-by-measure matrix has been obtained, the next step in the analysis involves assembling the measures into a "super-measure" of the desirability of each alternative. Utility theory provides the basis for the appropriate approach to aggregate the seemingly disparate measures. It is a logically consistent and tractable means of representing the degree to which each alternative fulfills decision maker's objectives. The use of utility theory ensures that any recommendation reflects:

- the interactions, if any, between measures
- the relative attractiveness of a specific level on a measure
- the relative attractiveness of performance on different measures.

For a more detailed presentation of these topics see Keeney and Raiffa (1976) and von Winterfeldt and Edwards (1986).

If the decision maker's preferences are consistent with some special independence conditions, then a multiple attribute utility model $u(x_1, x_2, \ldots, x_n)$, where x_i represents the level of performance on measure *i*, can be decomposed into an additive, multiplicative, or other well-structured forms that simplify assessment. An additive multiple attribute utility model can be represented as follows:

$$u(x_1, x_2, \ldots, x_n) = \sum_{i=1}^{n} w_i u_i(x_i) \quad (2)$$

where $u_i(\cdot)$ is a single attribute utility function over measure i that is scaled from 0 to 1, w_i is the weight for measure i and $\sum_{i=1}^{n} w_i = 1$.

If the decision maker's preference structure is not consistent with the additive model (2), then the following multiplicative model may be used, which is based on a weaker independence condition:

$$1 + ku(x_1, x_2, \ldots, x_n) = \prod_{i=1}^{n} [1 + kk_i u_i(x_i)] \quad (3)$$

where $u_i(\cdot)$ is also a single attribute utility function scaled from 0 to 1, the k_i's are positive scaling constants satisfying $0 \le k_i \le 1$, and k is an additional scaling constant that characterizes the interaction effect of different measures on preference. Methods for determining the value *k* can be found in Logical Decisions (1996), page 150. As a special case when $\sum_{i=1}^{n} k_i = 1$, the multiplicative model (3) reduces to the additive model (2). For a more detailed discussion of the assumptions underlying these two models, see Keeney and Raiffa (1976).

In this paper, we will assume that the decision maker's preference structure is independent and use the additive model in (2) in our analysis.

4 R&S EXPERIMENTAL SETUP

Assume that there are K ≥ 2 project configurations. For $2 \le i \le K$, let $\boldsymbol{X}_i = (X_{i1}, X_{i2}, \ldots, X_{in})$ denote a vector of random variables representing the performance measures for configuration *i*. Let $E[u(\boldsymbol{X}_i)]$ denote the expected utility (unknown) for configuration *i* and let

$$E[u(\boldsymbol{X}_{[1]})] \le E[u(\boldsymbol{X}_{[2]})] \le \ldots \le E[u(\boldsymbol{X}_{[K]})] \quad (4)$$

denote the ordered expected utility values. The goal is to select the project configuration with the largest expected utility $E[u(\boldsymbol{X}_{[K]})]$. If the R&S procedure achieves this goal a "correct selection" (CS) is made. The R&S procedure is designed to satisfy the following probability requirement:

$$P\{CS\} \ge P^* \text{ whenever } E[u(\boldsymbol{X}_{[K]})] - E[u(\boldsymbol{X}_{[K-1]})] \ge \delta^*$$

where $(1/K) < P^* < 1$ and $0 < \delta^* < 1$.

Figure 3 contains the flow of the combined analysis using MAU and R&S. The simulation model generates M (≥

1) replicates for each project configuration. On each replication for each configuration, the multiple attribute utility function in (2) is evaluated using the realization of $\boldsymbol{X}_i$. If the utility function realizations for a given configuration are not normally distributed then multiple replicates are conducted, over which the realizations are averaged. Then multiple replicates are made of the averages in order to produce approximately normal data for the R&S procedure. Goldsman et al. (1991) refer to this last step as making macroreplications. In our analysis, we used the two-stage indifference zone procedure for R&S due to Rinott (1978). In section 6, we illustrate this procedure on our project example. In the next section, we address the issue of selecting the indifference zone parameter δ^*.

5 SELECTION OF δ^*

In practice, the selection of δ^* depends on the judgment of the decision maker. Usually, it is assigned by appealing to the practical significance of the difference between the largest and second largest mean on a scale that has some physical meaning. For example, in the study conducted by Goldsman et al. (1991), four airline reservation systems were compared based on their expected time to failure. In this example, δ^* was selected on a time scale measuring system time to failure.

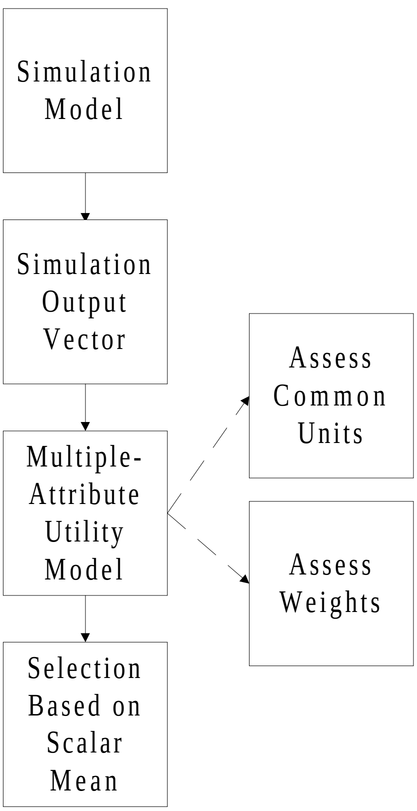

Figure 3: Analysis Flow

When R&S is based on expected utilities, the selection of δ^* can be challenging because δ^* has no direct physical meaning on the utility scale. To address this problem, we suggest a two-step process of first defining a δ_j^* for each j, $1 \le j \le n$, and then setting δ^* equal to

$$w_1\delta_1^* + w_2\,\delta_2^* + \ldots + w_n\,\delta_n^*.$$

In the first step, we use *certainty equivalents* defined on the single attribute utility functions. For a single attribute utility function, the certainty equivalent is equal to the inverse of the utility function evaluated at the expected utility (Clemen 1991, page 372). Let $\mathrm{E}[u_j(X_{ij})]$ be the expected utility for configuration i, $2 \le i \le \mathrm{K}$ on performance measure j, $1 \le j \le n$ and let

$$\mathrm{E}[u_j(\boldsymbol{X}_{[1]j})] \le \mathrm{E}[u_j(\boldsymbol{X}_{[2]j})] \le \ldots \le \mathrm{E}[u_j(\boldsymbol{X}_{[K]j})] \quad (5)$$

denote the ordered expected utility values. It is important to note that the ordering in (4) is not necessarily the same as the ordering in (5) because (4) depends on multiple performance measures. Let $CE_{[K]j}$ denote the certainty equivalent corresponding to $\mathrm{E}[u_j(\boldsymbol{X}_{[K]j})]$. Then, by definition,

$$\mathrm{E}[u_j(\boldsymbol{X}_{[K]j})] = u_j(CE_{[K]j}). \quad (6)$$

From (6), the quantity δ_j^* is defined by the R&S probability requirement:

$$P\{\mathrm{CS}\} \ge P^* \text{ whenever } u_j(CE_{[K]j}) - u_j(CE_{[K-1]j}) \ge \delta_j^*$$

where $(1/\mathrm{K}) < P^* < 1$ and $0 < \delta_j^* < 1$. To set δ_j^*, one can invert $u_j(CE_{[K]j})$, $u_j(CE_{[K-1]j})$ and establish an indifference zone based on $CE_{[K]j}$ and $CE_{[K-1]j}$. Since the latter two quantities are on the scale of the original performance measure, the decision maker should be able to establish an indifference zone more easily than on the utility scale. Once an indifference zone has been established on the scale of the original performance measure, the results can be substituted back into the utility function in order to establish δ_j^*.

It is important to note that for a constant indifference zone parameter value δ_j^* on the utility measure, the indifference zone on the original performance measure will be variable unless the utility function is linear. However, the indifference zone constructed on the performance measure axis need only be established for two points. If the utility curve accurately reflects the decision maker's preferences, then the zone defined by any other two points on the performance measure axis will adjust accordingly.

To demonstrate how the zone changes and to check for consistency, we find it helpful to plot an indifference-zone, preference-zone diagram (Bechhofer et al. 1995, page 178) for $CE_{[K]j}$ and $CE_{[K-1]j}$. The curve dividing the indifference-zone from the preference-zone is constructed by setting

$$u_j(CE_{[K]j}) - u_j(CE_{[K-1]j}) = \delta_j^*$$

and solving for $CE_{[K]j}$. For the utility function in (1), the resultant expression is

$$CE_{[K]j} = CE_{[K-1]j} - RT * Ln[-(\delta_j^* / B) * e^{(CE_{[K-1]j}/RT)} + 1]$$

Figure 4 contains three indifference-zone, preference-zone diagrams corresponding to δ_j^* equal to 0.2, 0.1, and 0.01, respectively, for the transport vehicle utilization example in Figure 2. The indifference-zone always appears below the curve and the preference-zone above as indicated on the graphs. As δ_j^* decreases the indifference zone parameter on the performance measure axis decreases and tends toward a constant value. This explains why the relationship becomes more linear as δ_j^* decreases.

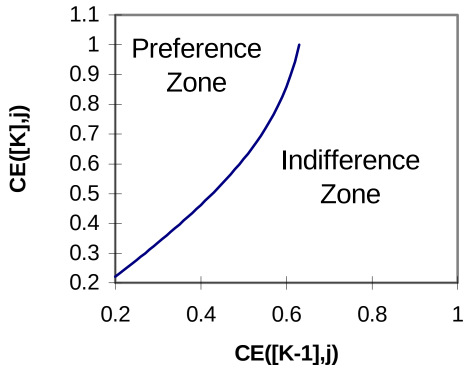

(a) For $\delta_j^* = 0.2$

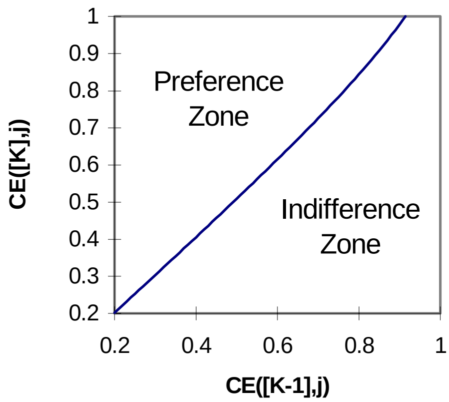

(b) For $\delta_j^* = 0.1$

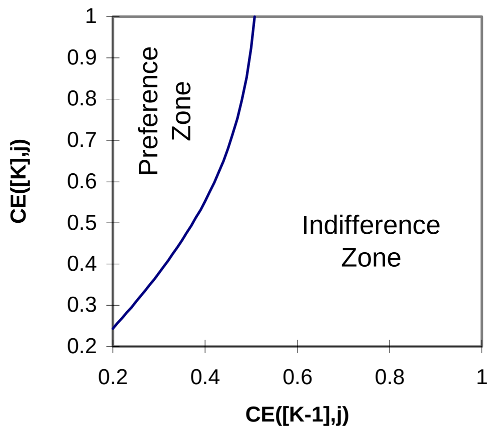

(c) For $\delta_j^* = 0.01$

Figure 4: Indifference-Zone, Preference-Zone Curves

We justify the second step of the procedure, i.e., to set δ^* equal to $w_1\delta_1^* + w_2\,\delta_2^* + \ldots + w_n\,\delta_n^*$ by noting that from (2) the E[u($\mathbf{X}_i$)] equals

$$w_1E[u_1(X_{i1})] + w_2E[u_2(X_{i2})] + \ldots + w_nE[u_n(X_{in})].$$

6 APPLICATION OF THE PROCEDURE

In this section, we illustrate our methodology on an example. Although the data used in the example are not real, they are representative. The simulation model generated results on a job that is realistic in both size and structure. Additionally, the utility functions and weights were assessed based on informal discussions with personnel who have field management experience.

We define the configurations based on resource levels along two dimensions: the number of source crews and the number of units of receiving equipment. Resource decisions along these two dimensions are considered the most important on a signal quality survey. We consider four configurations: one and two source crews in conjunction with 1100 and 1300 units of receiving equipment. All other resources and parameters remain fixed.

The performance measures include survey cost, survey duration and utilization for the following four crews: source, layout, transport, and packing. The recording crew is not included because it rarely bottlenecks production. The utility function for the survey cost and survey duration were defined over a range considered reasonable for a survey of the given size and complexity. Specifically, survey cost was defined over the range of 80 to 190 thousand US dollars with the following utility function:

$$1.004 - (7.52\text{E} - 05)e^{(X/2000)}$$

Job duration was defined over the range 240 to 480 hours with the following utility function:

$$1.002 - (6.16\text{E-}06)e^{(X/40)}.$$

The utility functions for the crew utilizations were all defined over the range 0.2 to 1.0 since a utilization of less that 0.2 would not be acceptable in the field. Utility functions for utilization are challenging to construct because management must balance two factors: desired productivity and worker satisfaction. To address this issue, we assigned two utility functions to each utilization performance measure, one for each factor. Both functions have the following form for all utilization measures:

$$1.019 - (2.769)e^{(-X/0.2)} \quad \text{(Desired Productivity)}$$

$$1 - (2.06\text{E-}09)e^{(X/0.05)} \quad \text{(Worker Satisfaction).}$$

There are at least two other ways to handle the balancing of these two factors. One way is to develop a single non-concave (or non-convex) function that increases over an initial range of utilization (since productivity outweighs worker dissatisfaction when workers are not overworked) and drops off when utilization gets too high (worker satisfaction outweighs productivity). We did not use this approach because a non-concave function causes technical difficulties for finding unique certainty equivalents. The second approach is to constrain the upper bound on the range of utilization to be something less than one. However, this approach forces the decision maker to provide a specific cut-off point beyond which a configuration would not even be considered.

The MAU function was constructed from a weighted sum of the ten single attribute utility functions. Weights were assigned as follows: cost (0.4), job duration (0.2), desired productivity for each utilization (0.05), and worker satisfaction for each utilization (0.05).

For the Rinott (1978) two stage procedure, we used the following parameters: $\delta^* = 0.0035$, $P^* = 0.9$. The δ^* was constructed using the technique described in Section 5. Recall that for the assessment of δ^*, any two points on the performance measure scale can be used. We chose to anchor the cost at $140 thousand and assessed an indifference zone value in the positive direction of $2,000. For job duration, our anchor point was 360 hours with an indifference zone value in the positive direction of two hours. All utilizations were anchored at 0.8 and assessed an indifference zone value of 0.01 in the positive direction.

In the first stage of the Rinott procedure, 100 simulation runs were made for each configuration. Since the MAU utility function values were not normally distributed, averages of the values were calculated based on batches of size ten yielding ten macroreplications for each configuration. The hypothesis of normality was not rejected for any of the samples based on the Chi-square, Kolmogorov-Smirnov, and Anderson-Darling tests in *BestFit* (Palisade, 1996).

Table 1 contains the results from the first stage of the R&S procedure. The configurations are numbered as follows: a single source crew with 1100 units of equipment (1), a single source crew with 1300 units of equipment (2), two source crews with 1100 units of equipment (3), and two source crews with 1300 units of equipment (4). The number of additional observations were calculated using the formulae and tables on pages 61-63 in Bechhofer et al. (1995). Note that each additional observation requires an additional ten simulation runs.

Table 1: First Stage Results from the R&S Procedure

Configuration	Average	St. Dev.	Add. Obs.
1	0.838	0.0052	10
2	0.960	0.0043	3
3	0.867	0.0044	4
4	0.925	0.0010	0

Table 2 contains second stage results. Configuration 2 is best since it has the highest sample average. These results reveal that an additional 200 units of equipment are more beneficial than an additional source crew. Adding both together is not worth the additional cost. A closer inspection of the data reveals that with 1100 units of equipment, equipment is the bottleneck. With 1300 hundred units, the transport vehicle crew becomes the bottleneck. Therefore, adding an additional source crew does not provide much additional benefit (compare configuration 1 and 3). In fact, the additional source crew becomes a detriment to cost and the utilization of the transport crew (close to 100 percent) when 1300 units of equipment are available (compare configurations 2 and 4).

Table 2: Second Stage Results from the R&S Procedure

Configuration	Average	St. Dev.
1	0.839	0.0069
2	0.960	0.0039
3	0.868	0.0039
4	0.925	0.0010

7 CONCLUSION

In this paper we have developed an R&S procedure applied to multiple project configurations that are evaluated on multiple performance measures. The core procedure relies on the ideas and techniques found in MAU theory. Our example demonstrates that it can be applied to realistic problems in which simulation is used.

We will consider three issues in future research: sensitivity analysis on the MAU assessments, other ways of assessing δ^*, and extending R&S methods to handle vectors of performance measure. The last issue would address the case in which it is difficult to make the assessments required for MAU theory.

ACKNOWLEDGEMENTS

The authors would like to thank Peter Highnam and Gary Lundeen of Schlumberger Austin Product Center for their help on and support of this work.

REFERENCES

Bechhofer, R. E., T. J. Santner, D. M. Goldsman. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. New York: John Wiley & Sons, Inc.

Clemen, R. T. 1991. *Making Hard Decisions*. Boston: PWS Kent Publishing.

Goldsman, D. M., B. L. Nelson, B. W. Schmeiser. 1991. Methods for Selecting the Best System. In *Proceedings of the 1991 Winter Simulation Conference*, ed. B. L. Nelson, W. D. Kelton, and G. M. Clark, 177-186. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Dyer, J. S. and H. W. Lorber. 1982. The Multi-attribute Evaluation of Program-Planning Contractors. *OMEGA*, 6:673-678.

Keeney, R. L. 1980. *Siting Energy Facilities*. New York: Wiley.

Keeney, R. L. and H. Raiffa. 1976. *Decisions with Multiple Objectives.* New York: Wiley.

Keeney, R. L. and D. von Winterfeldt. 1991. Eliciting Probabilities from Experts in Complex Technical Problems. *IEEE Transactions on Engineering Management* 38:191-201.

Keeney, R. L. and D. von Winterfeldt. 1994. Managing Nuclear Waste from Power Plants. *Risk Analysis* 14:107-130

Logical Decisions, Inc. 1996. *Logical Decisions for Windows,* Version 4.0. Golden, Colorado: Logical Decisions, Inc

Morrice, D. J., P. W. Mullarkey, A. S. Kenyon, H. Schwetman, J. Zhou. 1997. Simulation of a Signal Quality Survey. In *Proceedings of the 1997 Winter Simulation Conference,* ed. S. Andradottir, K. J. Healey, D. H. Withers, and B. L. Nelson, 1265-1272. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Palisade Corporation. 1996. *BestFit: Probability Distribution Fitting for Windows User's Guide.* Newfield, New York: Palisade Corporation.

Rinott, Y. 1978. On Two-Stage Procedures and Related Probability-Inequalities. *Communications of Statistics: Theory and Methods* A8:799-811.

Saaty, T. 1988. *The Analytical Hierarchy Process.* New York: McGraw-Hill, Inc.

Schoemaker, P. J. and C. C. Waid. 1982. An Experimental Comparison Of Different Approaches to Determining Weights in Additive Utility Models. *Management Science*, 28:182-196.

von Winterfeldt, D. and Edwards, W. 1986. *Decisional Analysis and Behavioral Research.* New York: Cambridge University Press.

AUTHOR BIOGRAPHIES

DOUGLAS J. MORRICE is an Associate Professor in the MSIS Department at The University of Texas at Austin. He received his undergraduate degree in Operations Research at Carleton University in Ottawa, Canada. He holds an M.S. and a Ph.D. in Operations Research and Industrial Engineering from Cornell University. Dr. Morrice has spent the last two years working on various projects at the Schlumberger Austin Product Center. His research interests include discrete event and qualitative simulation modeling and the statistical design and analysis of large scale simulation experiments. Dr. Morrice is a member INFORMS. He served as the Secretary for the INFORMS College on Simulation (1994-1996) and was Co-Editor of the *Proceedings of the 1996 Winter Simulation Conference*.

JOHN BUTLER is a PhD candidate in the College of Business Administration at the University of Texas at Austin. His research interests include decision theory, decision support systems, and mathematical modeling. Currently he is involved in a research project designed to aide in the selection of a technology for the disposition of US surplus weapons grade plutonium.

PETER W. MULLARKEY is a Research Scientist in the End-to-End Simulation Department at Schlumberger Austin Product Center - Research. He received his undergraduate degree in Civil Engineering at Worcester Polytechnic Institute in Worcester, Massachusetts. He holds a Ph.D. in Civil Engineering from Carnegie-Mellon University. Dr. Mullarkey has worked with Schlumberger doing research and development of knowledge-based decision support systems since 1985. Additionally, he worked on a two year project at Tandem Computers developing a state model-based diagnosis and maintenance system product.